IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Iyengar et al.	**Examiner: Sheng Jen Tsai**
Serial No: 10/715,225	**Group Art Unit: 2186**
Filed: November 17, 2003	**Docket: YOR920030488US1 (163-16)**

For: SYSTEM AND METHOD FOR ACHIEVING STRONG DATA CONSISTENCY

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Appeal from Group 2186

KEUSEY, TUTUNJIAN & BITETTO, P.C.
20 Crossways Park North, Suite 210
Woodbury, NY 11797
Tel: (516) 496-3868
Fax: (516) 496-3869

TABLE OF CONTENTS **Page(s)**

This Appeal is from a Final Office Action mailed on January 17, 2008 (referred to as the "Final Action") and an Advisory Action mailed on April 11, 2008 (Advisory Action). This Appeal was commenced by a Notice of Appeal and Pre-Appeal Brief Request for Review filed on April 17, 2008. A Notice of Panel Decision was mailed on June 9, 2008 indicating that the application remains under appeal, and Appellants hereby submit this Appeal Brief in furtherance of the Appeal.

I. REAL PARTY IN INTEREST

The real party in interest for the above-identified application is International Business Machines, Inc., the assignee of the entire right, title and interest in and to the subject application by virtue of an assignment of recorded in the U.S. Patent and Trademark Office.

II. RELATED APPEALS AND INTERFERENCES

There are no Appeals or Interferences known to Applicant, Applicant's representatives or the Assignee, which would directly affect or be indirectly affected by or have a bearing on the Board's decision in the pending Appeal.

III. STATUS OF CLAIMS

Claims 1~26 are pending, stand rejected and are under appeal. The claims are set forth in the attached Appendix.

IV. STATUS OF AMENDMENTS

No after final amendment was filed in this action.

V. SUMMARY OF CLAIMED SUBJECT MATTER

In general, the claimed inventions are directed to systems and methods for achieving data consistency among multiple copies. For purposes of illustration, the subject matter of the claims will be described with reference to certain Figures and corresponding text of Appellants' Specification (hereinafter, "Spec."), for example, but nothing herein shall be deemed as a limitation on the scope of the invention. For each Claim listed below, the claim elements are presented in italicized text, and are followed by citation to exemplary figures and/or supporting text in Appellants' Spec.

Claim 1 Recites:

In a system comprised of a plurality of storage elements, a method for maintaining objects in the storage elements comprising the steps of: **(see, e.g., FIG. 1, elements 12, 13)**

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements; **(see, e.g., FIG. 2, element 21; Spec. p. 11, lines 15-20; p. 18, lines 15-17; p. 19, lines 11-13).**

responding to a request to update an object by using maintained information to determine which of the storage elements may store a copy of the object; **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 18-19; p. 19, lines 14-19).**

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 19-21; p. 19, lines 18-19).**

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive. **(see, e.g., Spec. p. 17, lines 16-22; p.**

17. line 23 ~ p. 18, line 9; p. 18, lines 21-23; p. 19, lines 1-8; p. 19, lines 20-22).

Claim 10 Recites:

A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising: **(see, e.g., FIG. 1, elements 12, 13; p. 8, line 21 ~ p. 9, line 3).**

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements; **(see, e.g., FIG. 2, element 21; Spec. p. 11, lines 15-20; p. 18, lines 15-17; p. 19, lines 11-13).**

responding to a request to update an object by using maintained information to determine which of the storage elements may store a copy of the object; **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 18-19; p. 19, lines 14-19).**

*instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and***(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 19-21; p. 19, lines 18-19).**

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive. **(see, e.g., Spec. p. 17, lines 16-22; p. 17. line 23 ~ p. 18, line 9; p. 18, lines 21-23; p. 19, lines 1-8; p. 19, lines 20-22).**

Claim 11 Recites:

In a system comprised of a plurality of storage elements, a method for maintaining stored objects comprising the steps of: **(see, e.g., FIG. 1, elements 12, 13)**

maintaining a consistency coordinator which communicates with the storage elements

and stores information regarding which storage elements are storing which objects; **(see, e.g., FIG. 2, element 21; Spec. p. 11, lines 15-20; p. 18, lines 15-17; p. 19, lines 11-13).**

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object; **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 18-19; p. 19, lines 14-19).**

instructing each storage element in the set to invalidate a copy of the object; **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 19-21; p. 19, lines 18-19).**

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive. **(see, e.g., Spec. p. 17, lines 16-22; p. 17, line 23 ~ p. 18, line 9; p. 18, lines 21-23; p. 19, lines 1-8; p. 19, lines 20-22).**

Claim 17 recites:

A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising: **(see, e.g., FIG. 1, elements 12, 13; p. 8, line 21 ~ p. 9, line 3)**

maintaining a consistency coordinator which communicates with the storage elements and stores information regarding which storage elements are storing which objects; **(see, e.g., FIG. 2, element 21; Spec. p. 11, lines 15-20; p. 18, lines 15-17; p. 19, lines 11-13).**

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object; **(see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 18-19; p. 19, lines 14-19).**

instructing each storage element in the set to invalidate a copy of the object; and **(see,**

e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 19-21; p. 19, lines 18-19).

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive. **(see, e.g., Spec. p. 17, lines 16-22; p. 17. line 23 ~ p. 18, line 9; p. 18, lines 21-23; p. 19, lines 1-8; p. 19, lines 20-22).**

Claim 18 recites:

A system for maintaining strong data consistency comprising: **(see, e.g., FIG. 1, elements 12, 13)**

a plurality of storage elements; and **(see, e.g., FIG. 1, elements 12, 13)**

a consistency coordinator **(see, e.g., FIG. 1, elements 12, 13)**

configured to receive all updates for objects maintained in the storage elements, from object writers and content providers, the consistency coordinator communicating with the plurality of storage elements and maintaining information about which objects are stored in the plurality of storage elements, **(see, e.g., FIG. 2, element 21; Spec. p. 11, lines 15-20; p. 18, lines 15-17; p. 19, lines 11-13).**

*the consistency coordinator providing selective communication to storage elements which include an object to be updated such that for a given object update, the consistency coordinator instructs the storage elements that store a copy of the object to invalidate their copy of the object, (***see, e.g., Spec. p. 12, lines 4-9; p. 17, lines 8-15; p. 18, lines 18-19; p. 19, lines 14-19)** *and then delays updating of the object until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is unresponsive.* **((see, e.g., Spec. p. 17, lines 16-22; p. 17. line 23 ~ p. 18, line 9; p. 18, lines 21-23; p. 19, lines 1-8; p. 19, lines 20-22).**

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Claims 1-6, 10-12, 16-24 and 26 are rejected as being unpatentable over Iyengar (U.S. Pat. Pub. No. 2003/0172236) in view of Hiraoka et al.(U.S. Patent No. 4,733,348);

B. Claims 7-9, 14-15 and 25 are rejected as being unpatentable over Iyengar and Hiraoka and further in view of Chang (U.S. Pat. Pub. 2005/0128960).

VII. ARGUMENTS

A. The Combination of Iyengar and Hiraoka is *Legally Deficient* To Support a *Prima Facie* Case of Obviousness Against the Claimed Inventions

Appellants respectfully submit that the combination of Iyengar and Hiraoka is legally deficient to support a p*rima facie* case of obviousness against the subject matter of 1, 10, 11, 17 and 18. "In rejecting claims under 35 U.S.C. Section 103, the Examiner bears the initial burden of presenting a *prima facie* case of obviousness." *In re Rijckaert*, 9 F.3d 1531, 1532, 28 USPQ2d 1955, 1956 (Fed. Cir. 1993) (citing *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992)). Under 35 U.S.C. § 103, the factual inquiry into obviousness requires the Examiner to make a determination of: (1) the scope and content of the prior art; (2) the differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) any secondary considerations. *Graham v. John Deere Co.*, 383 U.S. 1, 17-18 (1966). "[Analysis [of whether the subject matter of a claim is obvious] need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." *KSR Int'l Co. v. Teleflex, Inc.*, 127 S. Ct. 1727, 1741 (2007). *See DyStar Textilfarben GmbH & Co. Deutschland KG v. CM. Patrick Co.*, 464 F.3d 1356, 1361 (Fed. Cir. 2006) ("The motivation need not be found in the references sought to be combined, but may be found in any

number of sources, including common knowledge, the prior art as a whole, or the nature of the problem itself."). The Examiner's analysis in support of the obviousness rejections should be made explicit. *KSR*, 127 S. Ct. at 1741.

In the case at bar, it is respectfully submitted that the Final Action fails, at the very least, to present a *prima facie* case of obviousness against claims 1, 10, 11, 17 and 18. Applicants respectfully assert that claims 1, 10, 11, 17 and 18 include features that are clearly not disclosed or suggested by Iyengar and Hiraoka, either singularly or in combination. More importantly, the Examiner's obviousness analysis is factually and legally flawed and based on unfounded interpretations of the claims as applied to the teachings of the cited references.

In general, the claimed inventions (claims 1, 10, 11, 17 and 18) are directed to systems and methods for managing objects to ensure cache consistency/coherency in a shared memory system where a plurality of caches may store a local copy of an object. A cache consistency coordinator is provided to receive *object update commands*, send instructions to those storage elements that are deemed to store a copy of the object (to be updated) to invalidate their local copy of the object, *wherein the updating of the object is delayed until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is unresponsive.*

1. Claim 1 is not Obvious in view of Iyengar and Hiraoka

In formulating the rejections, the Examiner seemingly relies on Iyengar as generally disclosing that a central cache may communicate with local caches to make sure that copies of an object to be updated are invalidated. However, Iyengar does not specifically teach that updating of the object is *delayed until it is determined that each storage element instructed to invalidate a*

copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive, as recited in the claimed inventions. In fact, the Examiner acknowledges at the very least that Iyengar does not specifically disclose *delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has . . . been deemed unresponsive.*

However, the Examiner seemingly relies on the teachings of Hiraoka in Col. 4, lines 41-50 as curing the deficiencies of Iyengar in this regard. But it is respectfully submitted that the Examiner's reliance on Hiraoka is wholly misplaced in this regard, as Hiraoka does not fairly teach or suggest *delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has . . . been deemed unresponsive.*

In the Advisory Action, the Examiner notes that with regard to the claimed feature of *delaying an updating of the object until it is determined (by the consistency coordinator) that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive,* the "core event" is "to invalidate a copy of the object. The Examiner further notes that the process by Hiraoka to purge a local copy of the TLB (table look aside buffer) is the same as "to invalidate a copy of the object." Indeed, the Examiner further notes in the Advisory Action that "a TLB certainly qualifies as an object and the purging the TLB would actually invalidate the TLB.

These arguments are *too* simplistic and are made out of context and with no due consideration of scope of the claimed inventions, as a whole. This can be readily seen by replacing the claim term "object" with TLB as follows: *delaying an updating of the* **TLB** *until it is determined that each storage element instructed to invalidate a copy of the* **TLB** *has either (i) acknowledged that it is not storing a valid copy of the* **TLB** *or (ii) been deemed unresponsive.*

This interpretation makes no sense on various levels. First of all, it is well known that a TLB (translation look aside buffer) is a cache of PTEs (page table entries), which is used by memory management hardware to improve the speed of virtual address translation. In general, most, if not all, state of the art processor technologies use a TLB, wherein a TLB has a fixed number of slots containing page table entries, which map virtual addresses onto physical addresses. In this regard, the Examiner's contention that the TLB is "certainly an object" that is stored in a cache is not factually correct, as the TLB is a cache structure that stores page table entries. In this regard, initializing or pursing a TLB is a process that involves purging the page table entries stored in the TLB cache structure.

Another fundamental flaw in the Examiner's analysis is that a "core event" (of the claimed process step in issue) is not *invalidating a copy of the object* as suggested by the Examiner, but rather *delaying an updating of the object* until certain events/conditions occur. In this regard, the Examiner's reliance on Hiraoka is misplaced. Even assuming that the **TLB** is a cached "object," there is nothing in Hiraoka that suggests *delaying an updating of the* **TLB** *until it is determined that each storage element instructed to invalidate a copy of the* **TLB** *has either (i) acknowledged that it is not storing a valid copy of the* **TLB** *or (ii) been deemed unresponsive*

The Examiner states (in the Advisory Action) that in Hiraoka, the "*source processor delays the updating* ... ", but the Examiner fails to explain how or where or what "updating a TLB" refers to in the context of the Hiraoka process. Indeed, even assuming, *arguendo*, that "purging" (initializing) a TLB can be deemed to be "invalidating a local copy of an object", in the proper context of the claimed invention, the Examiner simply ignores and fails to explain what would constitute "updating the TLB" in the Hiraoka system.

However, Hiraoka teaches that in a virtual memory control multiprocessor system, the purging (initialization) of TLBs of all process must be performed to equalize the contents of the TLBs (see Col. 1, lines 10-15). Hiraoka is merely directed to a system for purging the contents of TLB (table look aside butters) of a set of parallel processors. The teachings of Hiraoka regarding the process for purging TLBs of a multiprocessor system are very much different and irrelevant to the systems and methods as contemplated by the claimed inventions for maintaining consistency of copies of an object stored in storage elements.

In fact, as noted above, Hiraoka teaches that initialization (purging) of all processors must be completed at certain times for all processors. In this regard, there is nothing in Hiraoka that fairly teaches or suggests that the purging of one of the processors in a multiprocessor system can be skipped or disregarded (and subsequent tasks performed) in the event of a timeout condition (after issuance of a purge request signal) if one processor is unresponsive.

The Examiner relies on Col. 4, lines 41-50 of Hiraoka as teaching delaying updating until a processor is deemed unresponsive. However, Hiraoka teaches in Col. 4, lines 41-50 the following:

> The above operation can be performed when all the processors 20_0 through 20_3 are present. However, when the processor 20_3 is not present, the following operation is performed. The signal 48_3 representing that the processor 20_3 is not present is set at logic "1". The signal 48_3 of logic "1" is supplied to the OR gate 42_3. The OR gate 42_3 supplies the dummy TLB purge end signal to the AND gate 43. If the processor 20_3 is not present, the processor 20_0 can detect that all the TLB purge operations of the processors 20_0 through 20_2 are completed.

There is nothing in the cited paragraph relating to an unresponsive processor, but merely a condition in which the multiprocessor system comprise 3 processors instead of 4. As such, it is respectfully submitted that the Examiner's reliance on the cited passage is clearly misplaced.

In view of the above claim 1 includes features that are not disclosed or suggested by Iyengar and Hiraoka, either singularly or in combination. Accordingly, claim 1 is patentable over such combination.

2. Claim 10 is not Obvious in view of Iyengar and Hiraoka

It is respectfully submitted that the Final Action fails to set forth a factually and legally sufficient obviousness analysis, as required to support a *prima facie* case of obviousness against claim 10 based on the combination of Iyengar and Hiraoka. The obviousness rejection of claim 10 is based, essentially, on the same findings by the Examiner with regard to claim 1. However, for similar reasons discussed above with respect to claim 1, clearly, the cited combination does not teach or suggest, e.g., *wherein the updating of the object is delayed until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is unresponsive*, as essentially claimed in claim 10. Accordingly, claim 10 is patentable over such combination Iyengar and Hiraoka.

3. Claim 11 is not Obvious in view of Iyengar and Hiraoka

It is respectfully submitted that the Final Action fails to set forth a factually and legally sufficient obviousness analysis, as required to support a *prima facie* case of obviousness against claim 11 based on the combination of Iyengar and Hiraoka. The obviousness rejection of claim 11 is based, essentially, on the same findings by the Examiner with regard to claim 1. However, for similar reasons discussed above with respect to claim 1, clearly, the cited combination does not teach or suggest, e.g., *wherein the updating of the object is delayed until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is*

unresponsive, as essentially claimed in claim 11. Accordingly, claim 11 is patentable over such combination Iyengar and Hiraoka.

4. Claim 17 is not Obvious in view of Iyengar and Hiraoka

It is respectfully submitted that the Final Action fails to set forth a factually and legally sufficient obviousness analysis, as required to support a *prima facie* case of obviousness against claim 17 based on the combination of Iyengar and Hiraoka. The obviousness rejection of claim 17 is based, essentially, on the same findings by the Examiner with regard to claim 1. However, for similar reasons discussed above with respect to claim 1, clearly, the cited combination does not teach or suggest, e.g., *wherein the updating of the object is delayed until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is unresponsive*, as essentially claimed in claim 17. Accordingly, claim 17 is patentable over such combination Iyengar and Hiraoka.

5. Claim 18 is not Obvious in view of Iyengar and Hiraoka

It is respectfully submitted that the Final Action fails to set forth a factually and legally sufficient obviousness analysis, as required to support a *prima facie* case of obviousness against claim 18 based on the combination of Iyengar and Hiraoka. The obviousness rejection of claim 18 is based, essentially, on the same findings by the Examiner with regard to claim 1. However, for similar reasons discussed above with respect to claim 1, clearly, the cited combination does not teach or suggest, e.g., wherein the consistency coordinator *delays updating of the object until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii)*

is unresponsive., as essentially claimed in claim 18. Accordingly, claim 18 is patentable over such combination Iyengar and Hiraoka.

B. The Combination of Iyengar and Hiraoka and Chang is *Legally Deficient* To Support a *Prima Facie* Case of Obviousness Against the Claimed Inventions

With regard to the grounds of rejection listed as item B in Section VI above, rather than specifically address such rejection, it is suffice to say that the obviousness rejections of claims 7-9, 14-15 and 25 are legally deficient as a matter of fact and law at least for the same reasons given above for claim 1, 11, and 18 in view of Iyengar and Hiraoka and Chang at least by virtue of their dependence from claims 1, 11 or 18.

Accordingly, for at least the above reasons, it is respectfully requested that the Board reverse all claim rejections under 35 U.S.C. §103.

Respectfully submitted,

/Frank V. DeRosa/
Frank V. DeRosa
(Registration No. 43,584)

Mailing Address:

KEUSEY, TUTUNJIAN & BITETTO, P.C.
20 Crossways Park North, Suite 210
Woodbury, NY 11797
Tel: (516) 496-3868
Fax: (516) 496-3869

Claims Appendix

1. In a system comprised of a plurality of storage elements, a method for maintaining objects in the storage elements comprising the steps of:

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements;

responding to a request to update an object by using maintained information to determine which of the storage elements may store a copy of the object;

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

2. The method as recited in claim 1, wherein the step of maintaining information includes maintaining information regarding which storage elements are storing particular objects in the consistency coordinator.

3. The method as recited in claim 1, wherein the consistency coordinator includes multiple nodes and each node of the consistency coordinator stores information for a different set of objects.

4. The method as recited in claim 1, wherein the storage elements include at least one cache.

5. The method as recited in claim 1, wherein the storage elements are included in a distributed system.

6. The method as recited in claim 1, further comprising the step of obtaining a lock on the object to be updated before performing the update.

7. The method as recited in claim 1, further comprising the step of sending heart beat messages to obtain availability information about objects from the maintained information to a storage element and from a storage element to the maintained information.

8. The method as recited in claim 7, further comprising the step of declaring an entity down in response to failing to receive a heart beat.

9. The method as recited in claim 7, wherein the entity declares itself down in response to failing to receive a heart beat.

10. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising:

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements;

responding to a request to update an object by using maintained information to determine which of the storage elements may store a copy of the object;

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

11. In a system comprised of a plurality of storage elements, a method for maintaining stored objects comprising the steps of:

maintaining a consistency coordinator which communicates with the storage elements and stores information regarding which storage elements are storing which objects;

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object;

instructing each storage element in the set to invalidate a copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

12. The method as recited in claim 11, wherein the consistency coordinator includes multiple nodes and further comprising the step of at each node of the consistency coordinator, storing information about which storage elements are storing which objects for a different set of objects.

13. The method as recited in claim 11, further comprising obtaining a lock from the consistency coordinator by an entity attempting to update an object before performing the update.

14. The method as recited in claim 11, further comprising the step of sending, from the consistency coordinator to a storage element or from a storage element to the consistency coordinator, heart beat messages to obtain availability information.

15. The method as recited in claim 14, further comprising an entity expecting a heart beat, declaring itself down in response to failing to receive a heartbeat.

16. The method as recited in claim 11, wherein the storage elements include at least one cache.

17. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising:

maintaining a consistency coordinator which communicates with the storage elements and stores information regarding which storage elements are storing which objects;

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object;

instructing each storage element in the set to invalidate a copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

18. A system for maintaining strong data consistency comprising:

a plurality of storage elements; and

a consistency coordinator configured to receive all updates for objects maintained in the storage elements, from object writers and content providers, the consistency coordinator communicating with the plurality of storage elements and maintaining information about which objects are stored in the plurality of storage elements,

the consistency coordinator providing selective communication to storage elements which include an object to be updated such that for a given object update, the consistency coordinator instructs the storage elements that store a copy of the object to invalidate their copy of the object, and then delays updating of the object until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the objector or (ii) is unresponsive.

19. The system as recited in claim 18, further comprising a writer, which updates the object to be updated.

20. The system as recited in claim 19, wherein the writer resides on a same node as a storage element.

21. The system as recited in claim 19, wherein the writer writes an updated object to storage elements after the plurality of storage elements which are to receive the update have invalidated a current copy of the object.

22. The system as recited in claim 19, wherein the writer writes an updated object to storage elements after the plurality of storage elements which are to receive the update are determined to be unresponsive.

23. The system as recited in claim 18, further comprising at least one content provider.

24. The system as recited in claim 23, wherein the content provider resides on a same node as a storage element.

25. The system as recited in claim 18, further comprising heart beat messages, which may be transmitted between the consistency coordinator and the storage elements to obtain availability information from the consistency coordinator to a storage element or from a storage element to the consistency coordinator.

26. The system as recited in claim 18, wherein the storage elements include at least one cache.

Evidence Appendix

There is no evidence submitted pursuant to 37 CFR §§ 1.130, 1.131 or 1.132 or any other evidence entered by the examiner and relied upon by appellant in this Appeal.

Related Proceedings Appendix

None.